



Confluent Unveils Tableflow to Unite Analytics and Operations with Data Streaming

March 19, 2024

Tableflow makes it easier to convert streaming data to Apache Iceberg® tables to feed data warehouses, data lakes, and analytics engines

Fully managed connectors are enhanced with secure connectivity to private endpoints and a max 99.99% uptime SLA at a lower price

Stream Governance activates key data governance features automatically for all Confluent Cloud customers for safer sharing and scalability

LONDON--(BUSINESS WIRE)--Mar. 19, 2024-- [Confluent Inc.](#) (NASDAQ: CFLT), the data streaming pioneer, announced exciting new Confluent Cloud capabilities making it easier for customers to stream, connect, govern, and process data for more seamless experiences and timely insights while keeping their data safe. Confluent Tableflow easily transforms Apache Kafka® topics and the associated schemas to Apache Iceberg® tables with a single click to better supply data lakes and data warehouses. Confluent's fully managed connectors have been enhanced with new secure networking paths and up to [50 percent](#) lower throughput costs to enable more complete, safe, and cost-effective integrations. Stream Governance is now enabled by default across all regions with an improved SLA available for Schema Registry, making it easier to safely adjust and share data streams wherever they're being used.

For companies to make decisions that optimize costs, boost revenue, and drive innovation, it requires connecting the operational and analytical estates of data, which are traditionally siloed in organizations. The operational estate includes the SaaS applications, custom apps, and databases that power businesses such as Oracle, Salesforce, and ServiceNow. The analytical estate includes data warehouses, data lakes, and analytics engines that power analytics and decision-making and use data streams and historical tables to run queries and different analytical functions.

"The critical problem for modern companies is that operational and analytical estates must be highly connected, but are often built on point-to-point connections across dozens of tools," said Shaun Clowes, Chief Product Officer at Confluent. "Businesses are left with a spaghetti mess of data that is painful to navigate and starves the business of real-time insights."

Many organizations turn to Kafka as the standard for data streaming in the operational estate, and to [Iceberg](#) as the standard open table format for data sets in the analytical estate. Using Iceberg, companies can share data across teams and platforms while keeping tables updated as the data itself evolves.

Companies using Kafka want to utilize Iceberg to meet the rising demand for both streaming and batch-based analytics. As a result, many companies must execute complex migrations which can be resource-intensive, resulting in stale and untrustworthy data and increased costs.

"Open standards such as Apache Kafka and Apache Iceberg are popular choices for streaming data and managing data in tables for analytics engines," said Stewart Bond, Vice President of Data Intelligence and Integration Software at IDC. "However, there are still challenges for integrating real-time data across operational databases and analytics engines. Organizations should look for a solution that unifies the operational and analytical divide and manages the complexity of migrations, data formats, and schemas."

Tableflow makes it easier to feed data warehouses and data lakes for analytics

Tableflow, a new feature on Confluent Cloud, turns topics and schemas into Iceberg tables in one click to feed any data warehouse, data lake, or analytics engine for real-time or batch processing use cases. Tableflow works together with the existing capabilities of Confluent's data streaming platform, including Stream Governance features and stream processing with Apache Flink®, to unify the operational and analytical landscape.

Using Tableflow, customers can:

- **Make Kafka topics available as Iceberg tables in a single click**, along with any associated schemas
- **Ensure fresh, up-to-date Iceberg tables** are continuously updated with the latest streaming data from your enterprise and source systems
- **Deliver high-quality data products** by harnessing the power of the data streaming platform with Stream Governance and serverless Flink to clean, process, or enrich data in-stream so that only high-quality data products land in your data lake

Tableflow is currently available as part of an early access program and will soon be available for all Confluent Cloud customers.

More New Confluent Cloud Innovations

Connect brings new security, usability, and pricing enhancements to a portfolio of 80+ fully managed connectors

To build a central nervous system for a business, users have to be able to connect all of their data systems to capture continuous data streams. Connectors address the challenges of traditional data architectures that can silo data, decrease data quality, and lead to unplanned downtime. Connectors do this by seamlessly connecting data systems and applications as sources and sinks to Confluent Cloud. Confluent continues to add enhancements to connectors, a critical component of the data streaming platform, so that more users can experience fast, frictionless, and secure integrations.

With new upgrades to Connect, Confluent customers can:

- **Connect securely to critical data systems in private networks** using DNS Forwarding and Egress Access Points
- **Provision connectors reliably in seconds** with real-time configuration validations and a 99.99% uptime SLA
- **Stream data affordably at any scale** with up to 50% reduced data transfer costs of \$0.025/GB

Unlocking the full value of real-time data requires widespread connectivity with each of the data systems and applications running your business. Built together with our technology partners, the [Connect with Confluent \(CwC\)](#) partner program expands the data streaming ecosystem by providing easy access to fully managed data streams directly within the tools where teams are already working. This helps to simplify the development of real-time data products to share throughout the business.

Since CwC's launch last July with 17 technology partners, the program has seen massive growth with more than 40 partner integrations now included in the program. This quarter, CwC added new partners including [Advantco](#), [Aklivity](#), [Arroyo](#), [Asapio](#), [Census](#), [EMOX](#), [Kinetica](#), [Nstream](#), [Redis](#), [SingleStore](#), [Squid](#), and [Superblocks](#) —all having built new Confluent integrations within their applications.

Stream Governance improvements increase availability and reliability

Given today's increased focus on governance and compliance, Confluent is making it simpler for customers to take advantage of key Stream Governance features. Now all Confluent Cloud customers will have Stream Governance automatically enabled in their environments, providing easy access to key features including Schema Registry, [Data Portal](#), real-time Stream Lineage, and more, with support in all Confluent Cloud regions.

Schema Registry is a crucial component for governing data streams, helping teams enforce universal data standards to ensure data quality and data consistency while reducing operational complexity. The schemas stored in Schema Registry must be accessible to teams at all times, since any issues could lead to data compatibility errors and increased troubleshooting costs. To minimize these risks, [Stream Governance Advanced](#) now offers a 99.99% SLA for Schema Registry so organizations can avoid disruptions to critical workflows and manage compliance concerns. In addition, with Stream Governance features integrated into Tableflow, the benefits of improved uptime SLA for Schema Registry will soon be extended to Iceberg tables.

Enterprise clusters deliver more cost savings on more clouds

Confluent recently announced new [Enterprise clusters](#) with the same benefits as other Confluent Cloud clusters – including an industry-leading 99.99% SLA and a full ecosystem of enterprise-grade tools – plus enhanced security with private networking. This enables teams to uphold stringent security and networking requirements while simultaneously optimizing resources and cost-efficiency. Enterprise clusters automatically scale based on the workload with no manual intervention required. With new advancements in [Kora](#), Enterprise clusters can offer even more cost savings with a lower entry point and reduced throughput costs on both AWS and Microsoft Azure.

Learn more at Kafka Summit London 2024

Join the Confluent leadership team as they share their vision of streaming data products built by using Apache Kafka at the Kafka Summit London keynote on Tuesday, March 19, 2024, at 10:00 am GMT. [Register here](#) for free livestream access. A replay of the keynote presentation will be available [here](#).

No financial updates will be provided at the keynote. Confluent considered the product pricing changes announced in this release and at the keynote in its financial outlook for the first quarter of 2024 and fiscal year 2024, announced on February 7, 2024.

Additional resources

- Read the [Tableflow launch blog](#)
- Read the [Q1 '24 launch blog](#)
- Register for the [Q1 '24 Launch demo webinar](#)
- [Start your free trial of Confluent Cloud](#). No credit card required
- [Ask our Professional Services experts](#) for faster adoption

About Confluent

Confluent is the data streaming platform that is pioneering a fundamentally new category of data infrastructure that sets data in motion. Confluent's cloud-native offering is the foundational platform for data in motion – designed to be the intelligent connective tissue enabling real-time data, from multiple sources, to constantly stream across the organization. With Confluent, organizations can meet the new business imperative of delivering rich, digital front-end customer experiences and transitioning to sophisticated, real-time, software-driven backend operations. To learn more, please visit www.confluent.io.

Confluent and associated marks are trademarks or registered trademarks of Confluent, Inc.

This press release contains forward-looking statements, including among other things, statements regarding product pricing updates and our financial outlook, demand for Apache Iceberg, the benefits, performance, features and use cases of Tableflow, Connect, Stream Governance, and Enterprise clusters, and the adoption of Confluent's data streaming platform. The words "believe," "may," "will," "ahead," "estimate," "continue," "anticipate," "intend," "expect," "seek," "plan," "project," and similar expressions are intended to identify forward-looking statements. These forward-looking statements are subject to risks, uncertainties, and assumptions. If the risks materialize or assumptions prove incorrect, actual results could differ materially from the results implied by these forward-looking statements. Except as required by law, Confluent assumes no obligation to and does not currently intend to, update any such forward-looking statements after the date of this release.

Media Contact

Natalie Mangan

pr@confluent.io

Investor Contact

Shane Xie

investors@confluent.io

Source: Confluent, Inc.